009925_12Release.qxd 12/15/00 12:27 PM Page 1

ESTHER DYSON'S MONTHLY REPORT

Release 1.0®

VOLUME 18, NO. 11 | 19 DECEMBER 2000 | www.edventure.com

Peering into the World of Peer-to-Peer, Part II

INSIDE

PEERING INTO P2P, PART II 1
Peer and Now 2
Roots and vines
Is there a business here?
The search for standards
Distributed Computing 6
Popular Power
DataSynapse
United Devices
Peering Underneath the Hood 13
Virtual Net operating systems
OpenDesign
Box: Aimster
XDegrees
Grand Central
Box: Softwax
Distributed content distribution
The Next Stage 20
Resources & Contact Information 22
Calendar of High-Tech Events 23

Register now for PC Forum, March 25-28! Details at www.pcforum2001.com.

BY KEVIN WERBACH

Peer-to-peer (P2P) applications are buzzing about the Internet like swarms of bees. The record industry fears being stung by Napster's uncontrolled music sharing, while developers and investors envision a P2P revolution bringing the Net back to its decentralized origins . . . and forward to its pervasive future. Last month we surveyed the field of P2P applications. In this issue we look beyond the bees themselves to distributed computing, business dynamics and P2P infrastructure.

From today's custom-built, incompatible applications, the P2P world will become more defined. Application developers, service providers and infrastructure vendors will follow the paths best suited to their target markets and business models. P2P architectures will seep into the largest corporations in the form of distributed-computing services. Virtual operating systems will take hold. Will the P2P world maintain its open, anarchic, bleeding-edge enthusiasm as it matures? Who will develop the standards, formal or de facto, that enable interoperability? Which niches in the P2P ecosystem will be valuable, and which will turn into commodities? The adventure has just begun . . . .

Technological evolution is never uniform. Legacy code and user behavior change more slowly than the cutting edge. New applications and services often appear before back-end systems exist for them to realize their potential. In 1995, Amazon, eBay and Yahoo! were proofs-of-concept for the Internet as a commerce platform; only with massive infrastructure investment and shrewd management did they become significant revenue-generating businesses. Many of the technologies those companies now depend on, such as application servers, SSL security, content-management systems,

{ *continued on page 2* }

edventure
The conversation starts here.

009925_12Release.qxd 12/15/00 12:27 PM Page 2

load balancers, personalization engines, Javascript and content-delivery networks didn't exist when they were founded.

The P2P world is at that awkward stage when end-user applications are ahead of infrastructure and business models. Fortunately, many smart people are working to change that. While dozens of startups race to create the next killer P2P application, others are building platforms that support many uses. Some of these platforms manifest themselves as distributed-computing services for enterprises, while others show up as infrastructure or developer tools.

If there is a unifying theme in these P2P infrastructure developments, it's the idea of the network as a computer. (Apologies to Sun.) A well-functioning P2P service makes storing a file, processing a task or launching an application on another Internet connected machine as easy as on a single PC. By tapping intelligence at the edge and connecting machines to one another directly, the P2P approach hides the complexities and hierarchies inside the network. Because the Net doesn't really work like a computer, it must be hidden and abstracted. Eventually, the P2P vision means changing the Net itself . . . so that it *does* work like a computer. As Tim O'Reilly, founder of O'Reilly & Associates and an investor in Popular Power (SEE PAGE 8), explains, "We'll look at the Web as a big toolbox."

Peer and Now

Roots and vines

P2P applications, from Napster's music sharing to ICQ's instant messaging to Groove's collaboration toolkit (SEE **RELEASE 1.0**, NOVEMBER 2000), are like grape vines. Feed them water and nutrients, and they will produce succulent grapes. Throw in a skilled vintner and a bit of time, and you get a fine wine. The vines and grapes can't survive, however, without the subterranean root networks extending in every direction. Cut down the vines and more will spring up; poison the roots and the vines will die. Those who drink the wines won't understand the soil science and agronomy involved in keeping the roots healthy, but they appreciate the results in every glass.

Release 1.0® (ISSN 1047-935X) is published monthly except for a combined July/August issue by EDventure Holdings Inc., 104 Fifth Avenue, New York, NY 10011-6987; 1 (212) 924-8800; fax, 1 (212) 924-0240; www.edventure.com. It covers software, the Internet, e-commerce, convergence, online services, Web applications, mobile commerce, groupware, data networking, streaming media, enterprise applications, wireless communications, intellectual property and other unpredictable topics.

EDITOR-IN-CHIEF: Esther Dyson (edyson@edventure.com)

PUBLISHER: Daphne Kis (daphne@edventure.com)

EDITOR: Kevin Werbach (kevin@edventure.com)

SYSTEMS DIRECTOR: Scott Giering (scott@edventure.com)

MARKETING MANAGER: Joanna Douglas (joanna@edventure.com)

CIRCULATION MANAGER: Natasha Felshman (natasha@edventure.com)

OFFICE MANAGER: Bill Demeritt (bill@edventure.com)

EDITORIAL ASSISTANT: Kara Holmstrom (kara@edventure.com)

EXECUTIVE ASSISTANT: Irene Lawrence (irene@edventure.com)

Copyright © 2000, EDventure Holdings Inc. All rights reserved. No material in this publication may be reproduced without written permission; however, we gladly arrange for reprints, bulk orders or site licenses. Subscriptions cost $795 per year in the US, Canada and Mexico; $850 overseas.

009925_12Release.qxd 12/15/00 12:27 PM Page 3

Developers of early P2P applications had to grow the roots along with the vines. Napster, for example, includes its own desktop application, directory service, and file-storage system. Though these functions aren't unique to music sharing, Napster's Shawn Fanning had no choice but to build them. It's a familiar story. Early PC application developers had to create their own user-interface elements, but over time this and other functions migrated to the operating system or to widely available programming libraries and components. Today, with development environments such as Visual Basic, almost anyone can create a fully-functioning desktop application. Similarly, early Web developers had to write their own application platforms until commercial application servers arrived.

Many of the newer P2P platforms, such as Groove and Gnutella, already offer some degree of separation between the platform and the services on top. Groove, for example, allows developers to create specialized P2P applications without having to worry about functions such as message transport, security, firewall negotiation and data synchronization. Gnutella, as a collaborative project based on an open message protocol, encourages other developers to create specialized services that interoperate with its nodes.

The next step is for companies to focus purely on building the infrastructure for P2P services and applications. That's happening, along with other signs P2P is maturing from a phenomenon into a real force in the software and Internet industries.

Is there a business here?

P2P may work well as an architectural model, but as the foundation for new companies it must be joined with business models. As we've argued elsewhere (SEE RESOURCES SECTION) Napster is an anomaly, a valuable property simply because of its huge user base. ICQ and Hotmail were similar stories. Because of its legal troubles, though, Napster can't just sell itself to a large company that can monetize its traffic. In Napster's case, the obvious way forward is to create a subscription service for access to licensed, reliable music . . . which is exactly what the company has committed to do in partnership with Bertelsmann.

Others don't have the luxury of Napster's user base. From the beginning they must craft products and services that someone finds valuable enough to pay for. Of course, not all P2P technologies will be revenue-generating. Many will be open source or collaborative projects, such as Freenet and Gnutella, or hybrid open-source/commercial ventures, such as OpenCOLA, Jabber and Endeavors

009925_12Release.qxd 12/15/00 12:27 PM Page 4

Technology's Magi. As with Linux, participants in these efforts believe they can make money selling services or custom development while giving away the basic platform.

Commercial P2P efforts are pursuing three primary business models. The first is the Napster-like consumer subscription service. People pay for compact discs, they pay for cable television and they pay for AOL; it stands to reason some of them will pay for a P2P service that gives them something valuable. Part or all of the price may be subsidized by advertisers or sponsors. For music or other entertainment content, real-time data about consumer interests and opportunities to sell related products such as concert tickets are in many cases more valuable to the content producer than the product they give away.

The second major P2P business model is selling software or services to enterprises. P2P services promise lower costs, better scalability and greater flexibility than their client-server counterparts, all of which are characteristics that appeal to business customers. Assuming they can be convinced of its value, companies will pay for Groove for the same reason they bought Lotus Notes and other client-server groupware products. Most of the companies we discussed last month, such as Consilient, Roku and Infrasearch are taking this tack for at least part of their business. The distributed-computing startups discussed below (SEE PAGE 6) fit into this category as well, though primarily with a services rather than a software-licensing model.

Finally, there is the infrastructure approach. Website operators and enterprises pay companies such as Digital Island and Akamai because they make their services operate more efficiently. The infrastructure providers focus on hard technical and operational problems common to many sites, and are therefore able to offer a better solution than their customers could on their own. Or the infrastructure service is necessarily distributed across the network, as with content-delivery networks such as Akamai. Many of the companies described in this issue are taking the infrastructure tack, hoping to convince customers to buy rather than build.

To succeed, P2P companies will have to tailor their service offerings to the business model they choose. Enterprises have requirements around security, reliability and legacy integration. A service designed as invisible network infrastructure will look different than one marketed directly to users. P2P players must also have a view about which functions will be sources of proprietary value and which will be shared among all market participants . . . which leads us to a discussion of standards.

009925_12Release.qxd 12/15/00 12:27 PM Page 5

The search for standards

By their very nature, P2P applications and services depend on interconnection and network effects. The more machines that contribute their processing, storage and intelligence to the effort, the more value everyone receives. For that reason, network businesses drive toward common platforms, whether through market forces (e.g. Windows) or through formal standards (e.g. the Web and HTML).

This doesn't mean there will be one mega-standard for all things P2P. As Popular Power ceo Marc Hedlund articulates: "No one has yet made a compelling argument to me that the home Napster user and my enterprise pharmaceutical customer will mutually need to agree on an implementation language or application programming interface (API) for peer-to-peer." Within the distributed-computing realm, though, Hedlund sees a need for common abstraction layers and APIs. That way, developers can write applications that ignore the particulars of machines and service platforms. (SEE PAGE 13 FOR SOME ACTIVITY IN THIS DIRECTION.)

Technology standards usually take hold in one of three ways: a single player becomes dominant and everyone else ties into its platform (Windows, Palm, Oracle databases); enough big players find it in their interests to support a neutral standard rather than compete with proprietary solutions (wireless application protocol (WAP), extensible markup language (XML)); or a standard gains so much traction at the grass-roots developer level that the larger players must support it (HTML and most other basic Internet technologies, Linux, MP3). All of these strategies are now being tried in the P2P world, with limited success.

P2P seems ripe for the last of these scenarios, grassroots adoption, but there is little technical commonality among the various P2P efforts. The platforms with the greatest adoption and therefore the best positioning to become de facto standards, Napster and AOL's two instant messaging services, are among the least open to third parties. That may change. Napster has informally cooperated with others seeking to connect to its service by, for example, creating Macintosh client software before Napster itself did. Aimster (SEE PAGE 16) wants to turn AOL Instant Messenger buddy lists into the foundation of interoperable P2P applications, as does Jabber (SEE **RELEASE 1.0**, NOVEMBER 2000) with its open-source XML routing technology.

While this is going on, several of the biggest technology companies are pushing a standard called Universal Description Discovery and Integration (UDDI). Begun as a joint effort by IBM, Microsoft and Ariba in the realm of business-to-business (B2B) exchanges (SEE **RELEASE 1.0**, SEPTEMBER 1999), UDDI is a framework for P2P

009925_12Release.qxd 12/15/00 12:27 PM Page 6

services. According to IBM's Steve Burbeck, "UDDI is essentially Napster for generalized services." Unlike the grassroots efforts, UDDI was designed with the needs of large enterprises in mind, which may be a good or bad thing depending on your perspective. (The situation resembles the overlap between the RSS and ICE syndication standards; SEE **RELEASE 1.0**, JULY/AUGUST 1999.)

UDDI is a standards-based specification for companies to describe the service interfaces they offer (such as product catalogs), and to register them in a central directory. Other businesses can go to the directory, and based on the information stored there can plug their own systems directly into their partners' without the cumbersome point-to-point negotiations involved in traditional EDI connections. The standard uses XML and the simple object access protocol (SOAP), both open standards under the auspices of the World Wide Web Consortium, to format messages. The central registry makes UDDI not truly P2P, although the directory merely provides information that companies use to connect their systems together directly.

UDDI has been driven by large companies and particularly the technology vendors interested in selling to them. It serves a need, but only addresses part of the problem for companies seeking to integrate their services. Once remote services are discovered, the businesses involved still must have secure, reliable connections over which to created integrated services, and mechanisms to translate data between incompatible applications.

Thus it's not quite right to see UDDI as conflicting with bottom-up efforts to standardize P2P services. The two efforts simply have different world-views, despite the similarities in their goals. They may converge, battle, or thrive independently. Or another platform, such as Groove or one of the P2P infrastructure companies described below, may attract a critical mass of customers and developers.

Distributed Computing: "I Help Solve a Big Problem"

Distributed computing represents a special category of P2P infrastructure. These services turn computing into a utility. In a distributed-computing environment, a single computationally intensive or network-dispersed task, such as mining human-genome data to identify new drugs or determining Internet-wide network performance, is broken up into small chunks farmed out to a large number of machines. Typically, the client software on those machines includes a screen-saver, so that the

009925_12Release.qxd 12/15/00 12:27 PM Page 7

computation takes place only when the computer is otherwise idle, and pauses when a user types or clicks.

In our taxonomy of P2P applications, we placed distributed computing as a fourth major category alongside exchanges, collaboration and information discovery (SEE **RELEASE 1.0**, NOVEMBER 2000). Admittedly, though, distributed computing is the poorest fit with the P2P label. There is generally a central server that assigns the tasks and collates the results. Moreover, in such environments there is no direct communication between the peers, or if there is it is secondary. Still, all the actual processing is done independently on the client machines. Distributed computing efforts resemble other P2P areas in their use of untapped processing power at the edge of the network, though it's a stretch to call them peer-TO-peer.

Distributed computing is the most well-established branch of distributed systems, with research projects stretching back decades. Real-world examples were readily available before this year, such as SETI@Home and Distributed.Net. SETI@Home, conceived in 1997 and launched in May 1999, has over 2 million computers processing radio-telescope data to search for signs of intelligent extra-terrestrial life. Distributed.Net primarily focuses on cryptography, demonstrating that supposedly secure algorithms could be cracked through brute-force application of massive distributed processing power. Intel has used distributed-computing software for several years to farm out computationally difficult chip-design problems to idle workstations on its intranet.

To infinity and beyond

Recently, the Internet and the P2P movement have given distributed computing a shot in the arm. Platform Computing ceo Songnian Zhou, whose company sells resource-management software for distributed computing, did his PhD work in distributed computing in the 1980s. It's only today that commercial deployments are really taking off, he says: "After 15 years, we're suddenly waking up recognizing that distributed computing is going to be a big wave."

The massive well of untapped computing power connected to the Internet is an irresistible target. SETI@Home currently commands 28 teraflops of processing power, more than any physical supercomputer. Others are now putting the Internet-based distributed-computing model to work for more formal scientific research. The Grid Physics Network, announced in September and funded with a $12-million grant from the US National Science Foundation, will use distributed computing to process

009925_12Release.qxd 12/15/00 12:27 PM Page 8

the colossal data streams from experiments seeking to determine the structure of the universe. Sun and AOL have also announced plans to develop a distributed supercomputer using AOL's huge installed base of servers, though details are scant.

What's new this year is the proliferation of companies seeking to make distributed computing a for-profit business. Most envision a utility model, in which companies tap into distributed processing power on their own networks and can also buy capacity from a network of Internet users who contribute resources in return for some incentive (money, loyalty points, the satisfaction of helping to cure a disease). However, few have any real revenues to speak of, so it's premature to say whether the value proposition will pan out. All these companies face big challenges in assuring their corporate customers that security will be preserved.

And with so many competitors, it's easy to imagine computation becoming a commodity. Much of the academic work in network-based distributed processing failed to translate into commercial offerings because once the technology became stable enough, it made more sense for potential customers to build their own in-house parallel-processing server farms.

The companies discussed below think they can overcome these hurdles, and even if they don't, their experience will be an interesting test case for distributed computation in the real world. Popular Power, DataSynapse and United Devices are among the most interesting and substantial companies in this space, but they are not alone. Others include Applied Metacomputing, Centrata, Porivo, Parabon, Entropia and Distributed Science. In particular, Centrata, founded by MIT computer-science graduate students and funded by Kleiner Perkins, bears watching, though it won't publicly disclose its plans until the first half of 2001.

Popular Power: we the people

Popular Power's ceo, Marc Hedlund, previously ran the Internet division at Lucasfilm until the release of *Star Wars: Episode I*. In 1997, he pitched the idea of using distributed computing resources across the Internet to farm out computationally intense rendering tasks, but met with little interest in the company. A year later, Hedlund became interested in SETI@Home (ironically sponsored by Paramount, creator of Star Wars arch-rival Star Trek). He called his friend Nelson Minar, then at the MIT Media Lab, who had earlier done work on Internet-based distributed computing, and the two decided to to turn the concept into a business.

009925_12Release.qxd 12/15/00 12:27 PM Page 9

Calling around to large enterprises that were prospective customers, Hedlund found many already had efforts to tap the idle processing power on their internal networks. In most cases, though, the companies had concluded they didn't have the necessary expertise – or weren't willing to devote the necessary resources – to make the project work. Popular Power, like other distributed-computing startups, believes that by focusing specifically on the problem, it can deliver a solution where internal efforts failed.

POPULAR POWER INFO

Headquarters: San Francisco, CA
Founded: January 2000
Employees: 22
Funding: $1.6 million from individuals including Eric Schmidt (Novell), Ed Kozel (Cisco), Tim O'Reilly (O'Reilly & Associates), Brian Behlendorf (Collab.net) and Jolly Chen (Chen Capital)
URL: www.popularpower.com

Popular Power sees a few important factors distinguishing it from competitors. It is the first of the new breed of companies in the space to sign up customers and book revenues, says Hedlund. As its name implies, Popular Power sees its mission as people-centric rather than machine-oriented. Finally, where others have emphasized raw distributed computation, Popular Power sees the low-hanging fruit in distributed services that address network-wide issues in companies. "The people who were in the most pain were the people with network- or I/O-bound jobs," says Hedlund. As distributed processing power becomes more of a commodity, Popular Power believes these more network-aware applications will be a better business opportunity.

Examples include real-time network mapping, load testing, allocating bandwidth across multiple locations and devices, monitoring performance at all network nodes, managing distributed security or building distributed enterprise-wide databases. Companies such as Keynote sell network performance monitoring services and benchmarking reports, but they primarily collect data at large data centers run by companies such as Exodus. If Amazon.com wants to understand real-world network performance, argues Hedlund, it cares about latency to its customers, not to Exodus.

Popular Power's technical challenge is to make such services work reliably in aggregate despite the unreliability of their constituent nodes. Every Popular Power node has a "heartbeat" that tells the central dispatcher it's connected and operational, and the system tracks each node's reliability and assigns tasks accordingly. The parallels to an operating system kernel in an individual computer are strong, though there aren't yet robust abstraction layers through which distributed-computing companies can manage the network as cleanly as a developer or user does on a PC. Popular Power currently uses Java since the Java Virtual Machine is the farthest along in this regard, says Hedlund, but it's open to alternatives such as Microsoft's .Net framework should it mature sufficiently.

009925_12Release.qxd 12/15/00 12:27 PM Page 10

Popular Power is initially deploying is services on internal networks, although it is looking to a future that includes outsourcing over the Net. Hedlund uses a utility analogy: "My position is that I'm going to put electric outlets on all your walls and give you the best propane generator money can buy. Then I will go build a nuclear power plant, and when you're ready, you call."

Popular Power is using non-commercial projects such as optimizing influenza vaccination (in conjunction with the World Health Organization and Centers for Disease Control) to generate interest and build up a pool of contributors. The company's client software runs on Macintosh and Linux as well as Windows, making it one of the few multi-platform distributed-computing providers. Though Popular Power is currently angel-funded, Hedlund says the fact it has real customers such as BEA Systems has piqued VC interest.

DataSynapse: benefiting from broadband

DataSynapse targets the financial-services market, which isn't surprising when you discover that its founders founders came from Barclay's Global Investors and J.P. Morgan. Speed means everything in the financial world. Nanosecond advantages in execution or availability of information can mean billions. Traders want the best possible data and analytical models they can get . . . and they want them *now*.

Distributed processing promises to decrease costs, increase speed of simulations and other computationally difficult tasks and improve the quality of results traders depend on every day. Moreover, because financial markets are subject to intense demand surges, companies deploy capacity based on estimated peak loads. The ability to stockpile spare capacity and turn it up as needed is therefore appealing. "We see a wide range of companies today which have immediate need for a very flexible, industrial-strength distributed-computing platform," says DataSynapse ceo Peter Lee.

DataSynapse designed its software to support only broadband connections between its nodes. Other systems, following SETI@Home, take slow and intermittent dial-up connections into account, which means they look to get information uploaded as quickly as possible out of the distributed nodes into the central server. For long-duration projects such as searching for alien life, the resulting delay may be tolerable, but it becomes a problem in time-sensitive situations.

DATASYNAPSE INFO

Headquarters: New York, NY
Founded: March 2000
Employees: 33
Funding: $5.25 million from Rare Medium Ventures, NeoCarta Ventures, NYC Investment Fund, Silicon Alley Venture Partners, Wand Partners, and principals of Broadview International
URL: www.datasynapse.com

009925_12Release.qxd 12/15/00 12:27 PM Page 11

DataSynapse believes it's unique among distributed computing companies because it actually does use a P2P approach. "We do not use a hub-and-spoke distributed processing model," says Lee. DataSynapse nodes are deployed in a fully meshed topology, meaning they can talk to other nodes, can share intermediate results with one another, can break problems down into sub-problems for greater efficiency and can replicate data to eliminate bottlenecks. The company deploys its software on customer intranets and is also assembling a cadre of broadband-enabled members willing to donate processing capacity over the Internet. DataSynapse has partnered with fellow New York City startup Flooz.com (SEE **RELEASE 1.0**, OCTOBER 1999) to use Flooz as the incentive for members to contribute to the pool. (DISCLOSURE: KEVIN WERBACH'S WIFE IS AN EMPLOYEE OF FLOOZ.COM.)

The company licenses its technology to customers for internal deployment, and also charges for outsourced processing capacity from either its Internet-based members or partners it will plug in through an extranet. "Plugging into our network becomes as easy as plugging into the power grid," says Lee.

Lee has been encouraged that many potential customers already have considered or employed distributed computing. This means they will understand the concept, and he is confident DataSynapse can offer better results thanks to its P2P architecture. Security and reliability are obviously prime concerns in this market, so every communication is digitally signed, authenticated and encrypted. The system uses various measures to provide automatic failover and fault tolerance.

DataSynapse currently has seven beta customers in various stages of trials. The company officially launched last month, and plans to expand its marketing efforts significantly after the New Year.

United Devices: weaving the fabric of the Net

United Devices co-founder and ceo Ed Hubbard became interested in distributed computing while working at Dell in the office of cto Eric Harslem (a seed-round investor and advisor to United Devices). A participant in SETI@Home and distributed encryption-breaking efforts, he was convinced that a viable and profitable business could be built on the same kind of technology. Dell had SETI@Home nodes running on machines in its lab, but as a hardware vendor it wasn't the place for Hubbard to realize his vision. So he quit and started United Devices with Lance Hay, a former Intel auditor, who became cfo.

UNITED DEVICES INFO
Headquarters: Austin, TX
Founded: December 1999
Employees: 57
Funding: $13 million from SoftBank and Oak Investment Partners
URL: www.ud.com

009925_12Release.qxd 12/15/00 12:27 PM Page 12

Hubbard convinced David Anderson, the technical director of SETI@Home, to join United Devices as cto. More recently the company hired the 14 core developers from Distributed.Net, the largest distributed code-breaking effort, giving it expertise from the two most prominent non-profit distributed-computing projects. Hubbard argues this gives United Devices an important competitive advantage because its staff have the experience to avoid mistakes common among first-generation distributed-computing systems. "We've basically built a version-one product with 20/20 hindsight — a rare opportunity and one we're running with," he says.

United Devices has developed capacity-balancing technology to match the work assigned to a device to the capabilities of that device. A task that requires processing a 10-megabyte data file should be sent only to broadband-connected computers, for example, but when the files involved are small and can be processed offline for a significant length of time (as with SETI@Home), using narrowband-connected machines may be more efficient. Where DataSynapse trumpets its restriction to broadband networks, Hubbard sees United Devices' ability to handle all sorts of connections as evidence of its technical sophistication. Hubbard also emphasizes the lengths the company has gone to for security, including digitally signed code modules and Triple-DES encryption on all communications and stored data.

Though distributed processing seems destined to turn computation into a commodity, given the number of competitors in the space, Hubbard believes "it's not just about computation. Bandwidth, storage, geographic diversity, network diversity, device diversity, etc., all come into play." Hubbard sees United Devices' team giving it an edge here. He also believes the company's use of portable C will make it faster to connect customers' applications, most of which are written in C or C++, to the company's API than to those of Java-based competitors.

"I often refer to what we're doing as building an infrastructure to harness the 'fabric' of the Internet and corporate intranets," concludes Hubbard. That fabric can be applied to different tasks once companies such as United Devices create the abstraction layer on top of it. United Devices is offering services tapping both Internet- and intranet-based distributed computing. Each has different technical and operational requirements, though the core platform is the same.

Last month United Devices announced a deal with Exodus to put its distributed-processing network to work stress-testing Websites. The idea is that best way to simulate how sites will respond to millions of requests from all over the Internet is to send millions of requests from all over the Internet. The company is also working

009925_12Release.qxd 12/15/00 12:27 PM Page 13

with life-sciences companies including Incyte Genomics to speed bioinformatics projects. On the non-profit side, United Devices will power a molecular drug-design effort next year in conjunction with Oxford University and the National Foundation for Cancer Research that will use over 24 million hours of computing time.

Peering Underneath the Hood: P2P Infrastructure

Where distributed-computing companies use P2P technologies to deliver services directly to customers, others hope to build businesses focused on back-end P2P infrastructure. Similar to Internet infrastructure software companies such as Inktomi and Akamai, these vendors will make it easier to develop and deploy applications in a P2P world.

Virtual Net operating systems

Think about the functions a computer operating system (OS) performs. The OS interprets user requests and directs them to the application or device, identifies new peripherals as they are connected, manages memory and processor utilization between applications, offers an abstraction layer so that programmers need not write directly to each system component and presents users with an interface that hides the machine's internal complexity. A modern OS such as Windows, Linux, MacOS or Solaris employs central control mechanisms, but leaves most of the processing that users see to relatively autonomous applications and devices.

Without an OS, each new component added to a machine would have to be painstakingly integrated. Users and application developers would bear responsibility for allocating scarce resources of CPU, memory and disk space. Applications requiring storage would have to communicate directly with the hard disk controller chips, for example. Moreover, says OpenDesign's Edward Jung (SEE PAGE 14), "the PC OS abstraction allowed the benefits of hardware innovations to reach the application users. As buses, storage, and processing (video/CPU) became exponentially faster, the applications ran better."

The Internet today lacks an OS-like abstraction layer. "Internet applications are hard coded to the hardware, so the huge exponential increases in storage, bandwidth, and computation aren't as easily used by them without serious re-architecting," explains Jung. Developers of large Websites and network-scale applications face the Hercu-

009925_12Release.qxd 12/15/00 12:27 PM Page 14

lean tasks of scaling systems and integrating them with special-purpose subsystems such as storage networks and Web server farms. Businesses moving their processes online similarly must contend with the complexity of numerous point-to-point business relationships, most of which are manual. Connecting up heterogeneous systems over unreliable networks can easily become a nightmare.

As Shishir Mehrota, ceo of distributed-computing startup Centrata, puts it, "The goal is to build an application network that functions like physical networks, one in which you can be agnostic to individual pieces of hardware and the computing resources can be fluid." That's where the new generation of P2P network infrastructure players come in. Novell's Netware was an OS for managing controlled local-area networks, and Cisco's Internetwork Operating System (IOS) is an OS for the devices that route traffic around the Internet. Yet neither works at the higher levels of applications and content, especially across the open and heterogeneous Internet.

The hard part is coming up with an abstraction layer that works in such a complex and unreliable environment. "Everyone has grown up with a certain flavor of distributed computing called client-server, where you know the machine you're talking to," says Jung.

OpenDesign: scalable software architecture for the Web

OpenDesign, formerly known as WhatUWant, was spun out of Intellectual Ventures, a partnership of former Microsoft cto Nathan Myhrvold and ex-Microsoft chief software architect Edward Jung. Jung serves as OpenDesign's interim ceo, while Myhrvold is a major investor and advisor; the team also includes ex-Microsoft exec Steve Bush, Alex Cohen (who helped build the My Netscape portal and teaches in the film department at UC Berkeley in between stints at technology companies), as well as alumni of AOL, CNet and E-Trade.

OPENDESIGN INFO
Headquarters: Bellevue, WA
Founded: October 1999
Employees: 70
Funding: $7 million from founders
URL: www.opendesign.com

Given their experience with massive software-development efforts including Windows NT, E-Trade and Hotmail, OpenDesign's founders are intimately familiar with the pain of creating and scaling Internet-based applications and services. With OpenDesign, they hope to create a new network-wide abstraction layer to eliminate much of that pain.

"Today every Internet application has more computers than processes," explains Jung. Traditional PC software development assumed a single machine, owned by

009925_12Release.qxd 12/15/00 12:27 PM Page 15

one person or organization, running several concurrent applications and processes. In contrast, large Internet-based services such as MSN or CNet have a relatively limited number of core functions (authentication, directory lookups, Web serving, application serving, database access, email transport) each distributed across clusters of machines in data centers and often crossing administrative domains. Increasing PC performance is a simple matter of adding more processing or storage, but scaling a distributed Internet application is more challenging. Application load must be redistributed manually every time things change, and because adding servers (or new features) can create new conflicts.

The solution, Jung says, is to mimic the way the Internet and the Web work. In contrast to the software-development paradigm of integrate-then-deploy, Web content operates on a deploy-then-federate model. In other words, developers simply post their own Web pages and include hyperlinks that connect them to other pages. The Web's URL-based addressing system is smart enough to handle the behind-the-scenes interconnection between sites at the level of Web content. That addressing system sits on top of routers that dynamically balance network traffic loads through the magic of TCP/IP, augmented by caches and load balancers. Web developers need not worry about network topologies between their servers and their customers because the network itself is smart enough to do so. Once you get beyond static content and network usage levels scale beyond a certain point, however, the Net's built-in load balancing isn't good enough.

OpenDesign is creating a new platform that handles application-level load-balancing automatically. Using smart, programmable routers that can move and execute code, policies and data, the system changes network interconnections dynamically based on application demand, using both edge-node (P2P) and server-based processing as appropriate. Says Jung: "You can create an application-specific view of the network resources, and the system handles the network topology underneath to give you the view you want despite faults, load and other dynamic changes." On this platform, OpenDesign has built tools and services for functions such as directories, file management, messaging, user management, authentication, tracking, collaboration and policy enforcement.

For developers, OpenDesign's solution removes the need to build and manage expensive data-center hardware, and allows almost any programmer to create a load-balanced, fault-tolerant Internet application. For enterprises, OpenDesign will offer a management framework and also a link to its grassroots developer commu-

009925_12Release.qxd 12/15/00 12:27 PM Page 16

AIMSTER: INSTANT MESSAGING AND FILE SHARING TASTE GREAT TOGETHER!

The two most popular P2P services are Napster and instant messaging, so why not link them together? Aimster, a startup based at Rensselaer Polytechnic Institute's incubator facility in Troy, NY, has done exactly that. Aimster uses AOL Instant Messenger's namespace and transport network to link together users who can then share files with one another. The software has been downloaded over 1.5 million times, but ceo Johnny Deep has bigger plans.

Aimster is based on the online protocol (OP), similar to the Internet protocol (IP) but designed to connect people rather than machines. Where IP addressing ties users to specific physical locations on the Internet, OP links users based on presence. I'm still me, regardless of where I log in from. Deep also argues this presence-based addressing structure promotes privacy, because individuals aren't traceable to physical locations.

Deep hopes to become the addressing infrastructure for other P2P services such as Napster, which up to now have used their own proprietary namespaces. Aimster is also close to releasing a new version of its software that offers an integrated desktop and Web browser in addition to the current messaging and file-sharing services. With such a P2P browser (Deep calls it "the Browster"), users could sign up to community-oriented buddy lists and receive information that matched their interests. Aimster has done one project involving the band Radiohead, which used the service to promote its recent album *Kid A*. 37,000 AOL users signed up for the Radiohead buddy list in 32 hours.

nity. Customers can connect in legacy code, but to get the full benefits of the platform will require some re-writing of existing applications.

OpenDesign is engaged in pilots and plans to release a public beta in the first quarter, with full commercial availability of the SDK by mid-2001. It will charge a mix of licensing fees for the base platform and recurring fees for additional services.

XDegrees: turning up the heat

"The past is littered with failed attempts to use the Web as application architecture, when it's clearly designed best to serve as a publishing medium. A distributed architecture intelligently pushes applications out across the network, and may or may not use the Web as an interface to those applications," says XDegrees ceo Michael Tanne. The founder of ad server vendor AdForce, which was acquired by CMGI, Tanne did a stint as entrepreneur-in-residence at Redpoint before joining with several founders of Junglee to put together XDegrees.

XDEGREES INFO

Headquarters: Mountain View, CA
Founded: July 2000
Employees: 15
Funding: $8 million from Redpoint and Cambrian Ventures
URL: www.xdegrees.com

Just as AdForce and DoubleClick quickly superseded in-house ad rotation scripts in most companies because they built a better solution by making it their core competency, Tanne says, P2P companies will be better off outsourcing infrastructure to specialists such as XDegrees. "An application provider should not be required to build the whole stack," he says, because doing so involves higher costs and leads to fragmentation among competing vertical silos.

009925_12Release.qxd 12/15/00 12:27 PM Page 17

Like Groove Networks (SEE **RELEASE 1.0**, NOVEMBER 2000), XDegrees will offer application developers a platform that handles low-level functions transparently, allowing those developers to focus on their unique applications. Unlike Groove, though, XDegrees isn't associating itself with a specific class of applications such as collaboration and groupware. Tanne sees network-wide applications such as virus checking, spam blocking, collaboration and caching, as well as distributed versions of traditional enterprise applications such as sales-force automation as the most promising areas for P2P development. These services will function best, Tanne argues, when combining elements of P2P and client-server architectures as needed.

Recognizing the importance of addressing and people in a P2P environment, XDegrees has developed a naming system that includes people as a class of object along with data, devices and applications. According to Tanne, "XDegrees' naming infrastructure creates an abstraction between logical references and physical resources for a file, an application, a user, etc." In other words, individuals can be referenced intelligently regardless of their location or the device they use to connect to the network. The naming system also confirms that all copies of the same document are logically equivalent. In other words, content creators can manage documents on the network, and users can be assured the copies they see are intact, while still allowing the files to be distributed around the network in P2P fashion. Billing and digital rights management will also be available tied to the naming system.

XDegrees' system integrates intelligent caching to ensure that resources are available regardless of network conditions or whether an individual user is connected. It also builds in storage and a scalable messaging backplane capable of handling hundreds of millions of nodes and billions of accesses per day. To encourage developer adoption, existing applications can be integrated with XDegrees platform using established standards such as HTTP, XML and SOAP. The company will also offer add-in modules allowing developers to drop in functions such as file sharing, download management and user management, with the goal of further speeding adoption.

The company plans to offer its infrastructure as a hosted service, charging a combination of up-front licensing and usage-based fees to developers.

Grand Central: making the trains run on time

CNet founder Halsey Minor says that when he came up with the idea for Grand Central two years ago, people thought he was crazy. So after leaving CNet he created a vehicle, Twelve Entrepreneuring, to develop business ideas he and partners such as

009925_12Release.qxd 12/15/00 12:27 PM Page 18

SOFTWAX: CUSTOM P2P SOLUTIONS

Softwax, a 16-person startup founded in September 1999 and located in New York City, develops P2P solutions for corporations, online communities and media producers. The company was started by ceo Jason Grosfeld, who previously worked for the hedge fund that made the first investment in Napster. President Matt Epstein was an attorney at Latham & Watkins, and cto David Rosenbloom joined from Web development and consulting firm Rare Medium, where he held the same title.

The company has created a set of modular components to handle functions such as storage, security, metadata, search and authentication, which it can customize into P2P applications. It sees market opportunities in creating secure P2P file-sharing applications for media companies, offering community-oriented Websites new tools for member communication and giving companies new ways to collaborate and share documents.

Xerox PARC chief scientist John Seely Brown and *NetGain* author John Hagel hatched. Grand Central, which is Twelve's first company, plans to begin beta testing its service in February.

"Five or ten years from now, all global businesses will look like CNet and Yahoo!" says Minor. These companies sit at the intersection of thousands of data-sharing relationships, and their core competency is in aggregating and managing that information to create value. In Minor's formulation, "their business is the sum of those relationships." Where the first generation of Internet companies were largely based on person-to-machine interaction, these new virtual businesses will depend on automated machine-to-machine interactions to move information across application and company boundaries. Yet the Internet wasn't built with a messaging backplane to facilitate, manage, track and bill for such data-sharing relationships in a secure, robust, interoperable way. That's where Grand Central comes in.

Grand Central will operate a virtual overlay network of servers that companies can connect to in order to exchange business information. Users are authenticated at the edge of the Grand Central network and can then interact securely with all other users. They can also create routing rules to connect applications into network-spanning services, much as Unix allows data to be piped through a series of applications. For example, a company could rout information from its customer-relationship management system into its general-ledger application and through a third-party credit-scoring service. Grand Central won't offer any applications itself, but will provide a software development kit so customers can plug their own applications into the network.

GRAND CENTRAL INFO

Headquarters: San Francisco, CA
Founded: May 2000
Employees: undisclosed
Funding: Incubated by Twelve Entrepreneuring
URL: www.grandcentral.com

In addition to security and robust messaging, Grand Central will offer a directory for resource discovery and an integrated billing system. The directory element resembles the vision of the UDDI group (SEE PAGE 5), but discovery alone doesn't address the core problem of

009925_12Release.qxd 12/15/00 12:27 PM Page 19

establishing relationships and data flows between multiple providers. Minor says that compared to existing server-based data-sharing solutions such as Vitria and WebMethods, Grand Central's network will be far cheaper to deploy and better able to support the complexity of thousands of inter-connected and evolving relationships. The company will charge on a tiered usage basis, with an initial level of activity free to encourage developers onto the platform.

Grand Central isn't entirely P2P because it uses servers and central directories, but it represents a more distributed alternative to middleware bridges that route data through central control points. In Minor's vision, information flows through Grand Central much as it flows through the Internet, without defined meet points in the middle. The primary difference is the higher level of functionality to support scalable business relationships.

Distributed content distribution

The P2P services we have described in this section so far are invoked by companies seeking to share data and resources with one another. There are also network-based services that stand on their own. Just as Napster users benefit from the network effect of many participants sharing music files, customers desiring virus-checking or spam-filtering may benefit from information drawn from others on the network. A suspicious file on one machine may be nothing, but the same suspicious file on twenty or a thousand machines suggests that something is afoot . . . and a network-based service can use that information to protect others.

We described several such "network-scale services" last month, including search and information discovery. Content distribution is an example of a network-scale service that can operate within the Internet's infrastructure. Content-delivery networks (CDNs) such as Akamai and Digital Island push content to the edge of the network, improving performance (SEE **RELEASE 1.0**, DECEMBER 1999). For the most part, though, CDNs are centrally controlled, using hierarchical caches or "air-traffic control" centers to manage the flow of information. If the point is to move processing and storage to the edge, why not do so in a P2P configuration?

The centralized CDN model works best for large publishers such as Yahoo! and CNN that broadcast content to millions of users. Yet for situations where information is coming from many sources, and usage patterns are changing rapidly, this approach faces scaling issues. Fortunately, says Kalepa ceo and former Sun Java evangelist Miko Matsumura, "the wonderful good news is that on the edge of the net-

009925_12Release.qxd 12/15/00 12:27 PM Page 20

work there's an increasing proliferation of what I would consider to be server-class machines." And those machines can be tapped to distribute content more efficiently. Some companies already use P2P mechanisms for distributing files. MyCIO, a subsidiary of Network Associates, announced last month its Rumor technology for virus updates. Each Rumor-enabled client checks first whether other machines on the local network have new virus definitions, instead of going over the Net to a central server each time. This cuts down on network load and allows patches to spread more quickly. 2AM, an online games developer, uses a similar mechanism in its Streaming Peer platform for distributing rich media files to end-users.

P2P content delivery will soon become more prominent. MangoSoft sells a product that ties together the browser caches of each user on a local-area network. Some end-user P2P platforms, including MojoNation and OpenCOLA (SEE **RELEASE 1.0**, NOVEMBER 2000), include protocols for distributed delivery of large objects across the network. In addition, several stealth-mode startups, including eMikolo and Kalepa, are developing peer-based content-delivery infrastructure services. eMikolo, based in New York with R&D in Israel, recently announced $4.5 million in funding from Israel Seed Partners and individuals such as former Excite@Home svp Craig Donato. Kalepa, which has attracted attention because of Matsumura's high-profile role at Sun, has raised $1 million led by Angel Investors and Jump Investors.

Neither will talk yet, though Matsumura does venture obliquely that "it's going to be important for there to be players that don't focus necessarily on the outer edge of the network, but more on the inner edge of the network."

The Next Stage

Ironically, the collapse of Internet stocks is occurring just as the Net turns the corner in three significant ways. We shouldn't be surprised. The PC industry lost its early exuberance in the late 1980s and early 1990s, resulting in the death or decline of many companies. But it was at that time that PCs became part of mainstream life in the US and many other countries, and that a new generation of technology vendors such as Microsoft, Compaq and Dell rose to dominate the industry. In the Internet world, major developments bear watching.

First, companies are making significant investments in Internet-based technologies. General Motors putting up a corporate Website was an important symbol that the

009925_12Release.qxd 12/15/00 12:27 PM Page 21

Internet was real; GM re-engineering its supply chain around Internet-based technologies means far more in real terms. The tens of billions of dollars in annual spending involved is only half the story; such investments reorient thousands of jobs within such companies and their partners, change performance expectations in radiating circles out through interlinked global partner networks and set the roadmap for Internet-based software and networking companies.

The pattern was the same in the early days of the PC era, the previous computing revolution. When Wall Street financial analysts snuck Apple IIs in the back door so they could run VisiCalc it was thrilling, but the big and lasting changes occurred when businesses began running their operations on Lotus 123 or its successors and putting PCs on every desktop.

The second great Internet milestone is the growth in broadband and non-PC connections. As we have discussed before (SEE **RELEASE 1.0**, FEBRUARY 1999), always-on high-speed connections promise a different and more influential Internet experience than most users see today. In parallel, the expansion of connectivity to non-PC devices such as mobile phones, Sony Playstation game consoles and wireless personal digital assistants such as the Palm VII is taking the Net outside the narrow realm of computing and turning into a constant presence in daily life.

The third Internet transformation is the one we have been discussing in the past two issues. The business activity of the first milestone, coupled with the individual usage levels of the second, demand Internet services of a scale earlier architectures cannot achieve. The Internet of the future must operate in an environment of unprecedented heterogeneity, must do so robustly and must make it as simple to deploy a network-scale application as it is today to create a static Web page.

The PC never would have reached its current level of influence without graphical user interfaces that made using and developing for them easy and familiar. By the same token, the Net won't achieve its manifest destiny without a new kind of infrastructure. As we have discussed, that infrastructure must be radically distributed in order to scale, and it must offer abstraction layers on top of distributed heterogeneous hardware and networks that hide complexity and take advantage of the wealth of resources throughout the network. For these reasons, P2P will be important to the future of the Net, and the foundation of many successful businesses. ■ R 1.0

COMING SOON

- Broadband service platforms.
- How safe is your site?
- User experience.
- Triumph of the Weblogs.

- And much more. . . (If you know of any good examples of the categories listed above, please let us know.)

009925_12Release.qxd 12/15/00 12:27 PM Page 22

Resources & Contact Information

Clay Shirky, Saul Klein, **Accelerator Group**, 1 (212) 274-1305; fax, 1 (212) 275-1848; clay@shirky.com, saul@acceleratorgroup.com

John Deep, **Aimster**, 1 (518) 229-3319; johnd@aimster.com

Shishir Mehrota, **Centrata**, 1 (650) 561-0222; fax, 1 (650) 561-0201; shishir@centrata.com

Peter Lee, Mark Roth, **DataSynapse**, 1 (212) 842-8842; fax, 1 (212) 842-8843, peter@datasynapse.com, mark@datasynapse.com

Halsey Minor, **Grand Central/Twelve Entrepreneuring**, 1 (415) 277-9400; fax, 1 (415) 277-9401; halsey@twelve.com

Steve Burbeck, **IBM**, 1 (919) 254-0732; sburbeck@us.ibm.com

Miko Matsumura , **Kalepa**, 1 (650) 329-8888; miko@kalepa.com

Alex Cohen, **OpenDesign**, 1 (415) 269-0557; xcohen@opendesign.com

Edward Jung, **OpenDesign**, 1 (425) 467-5510; ed@opendesign.com

Tim O'Reilly, **O'Reilly & Associates**, 1 (707) 829-0515, tim@oreilly.com

Songnian Zhou, **Platform Computing**, 1 (905) 948-4200; zhou@platform.com

Marc Hedlund, Nelson Minar, **Popular Power**, 1 (415) 402-0683; marc@popularpower.com, nelson@popularpower.com

Jason Grosfeld, Matthew Epstein, David Rosenbloom, **Softwax**, 1 (212) 242-2107; fax, 1 (212) 202-3966; jason@softwax.com, matt@softwax.com

Ed Hubbard, **United Devices**, 1 (512) 630-6490; fax, 1 (512) 331-6235; ed@ud.com

Michael Tanne, **XDegrees**, 1 (650) 691-0400; fax, 1 (650) 691-0490; mtanne@xdegrees.com

For further reading:

O'Reilly & Associates P2P site, www.openp2p.com

"Peer-to-Peer Architectures and the Magi Open-Source Infrastructure," www.endtech.com/html/Architecture.html

Kevin Werbach, "Facing the Music: Beyond the Napster Injunction" (originally published in the *Industry Standard*): release1.edventure.com/release3.cfm?Counter=2830760

009925_12Release.qxd 12/15/00 12:27 PM Page 23

Calendar of High-Tech Events

2001

JANUARY 6-9	**CONSUMER ELECTRONICS SHOW 2001** - Las Vegas, NV. The world's largest annual tradeshow for consumer technology. Keynote speakers include Bill Gates and Craig Barrett. To register, call 1 (888) 237-7469; fax, 1 (703) 907-7675; email, CESinfo@CE.org; www.cesweb.org.
JANUARY 22-24	**SHOWCASE 2001** - Palm Springs, CA. Presented by Upside Media. It's time for the showcase showdown! For more info, call 1 (415) 908-2650; fax, 1 (415) 908-2660; www.upsideevents.com/showcase.
JANUARY 25-30	**WORLD ECONOMIC FORUM** - Davos, Switzerland. This year's Forum will once again bring together top business and political leaders, concerned citizens and creative academic thinkers. For info, see www.weforum.org. E
JANUARY 29-FEBRUARY 1	**COMNET** - Washington, DC. Network in the new year and hear keynotes by Colin Powell, John Roth and Dennis Strigl. Register by fax, 1 (508) 759-4552; www.comnetexpo.com.
JANUARY 30-FEBRUARY 2	**LINUXWORLD CONFERENCE & EXPO** - New York, NY. Learn to develop, integrate, and run a business on the Linux OS, with general championing of Linux along the way. For more info, call 1 (800) 657-1474; email, linuxworldexpo@idg.com; www.linuxworldexpo.com/papers.html.
FEBRUARY 4-6	**DEMO 2001** - Phoenix, AZ. Chris Shipley picks the hot startups again. For more info, call 1 (650) 312-0545; fax, 1 (650) 286-2750; email, registrar@demo.com; www.demo.com.
FEBRUARY 10-14	**MILIA** - Cannes, France. For professionals involved in the creation and distribution of interactive content. For info, contact Ted Baracos at +33 (1) 41 90 44 80; fax, +33 (1) 41 90 44 70; email, info@milia.com; www.milia.com.
FEBRUARY 14-16	**O'REILLY PEER-TO-PEER CONFERENCE** - San Francisco, CA. Gee, maybe there's something to this peer-to-peer thing? Meet some of the companies covered in this issue. For information, visit conferences.oreilly.com/p2p/.
FEBRUARY 21-24	**TED 11** - Monterey, CA. Richard Saul Wurman invites you to celebrate youth, age, the two words "will" and "still" and presentations by people under 30 and over 70. For more info, call 1 (401) 848-2299; fax, 1 (401) 848-2599; email, wurman@ted.com; www.ted.com.
MARCH 25-28	**PC FORUM** - Scottsdale, AZ. Join us! Registration is now open at www.pcforum2001.com. For more information contact Daphne Kis at 1 (212) 924-8800; daphne@edventure.com. E K

E Events Esther plans to attend.
K Events Kevin plans to attend.

Lack of a symbol is no indication of lack of merit. The full, current calendar is available on our Website, www.edventure.com. Please contact Kara Holmstrom (kara@edventure.com) to let us know about other events we should include.

009925_12Release.qxd 12/15/00 12:27 PM Page 24

Register now for PC Forum! This year's theme is "Define Yourself." Registration forms and additional information are available at **http://www.pcforum2001.com**.

Join our email list! We've started a free email newsletter, *The Conversation Continues*, for commentary, industry analysis and pointers to interesting Websites. To sign up, please visit **http://release1.edventure.com/conversation**, or send email to **conversation@edventure.com** and you'll automatically be added to the list.

Do we have your email address? *Release 1.0* subscribers can now download each month's issue electronically. If you have not already done so, please send your email address to natasha@edventure.com or fax it to 1 (212) 924-0240 in order to enable online access.

Release 1.0 Subscription Form

Complete this form and join the other industry executives who regularly rely on *Release 1.0* to stay ahead of the headlines. Or if you wish, you can also subscribe online at www.release1-0.com.

Your annual *Release 1.0* subscription costs **$795** per year (**$850** outside the US, Canada and Mexico), and includes both the print and electronic versions of 11 monthly issues; 25% off the cover price when you order from our online archives; a *Release 1.0* binder; the bound transcript of this year's PC Forum (a $300 value) and an invitation to next year's PC Forum.

NAME _____

TITLE _____ COMPANY _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____ COUNTRY _____

TELEPHONE _____ FAX _____

EMAIL* _____ URL _____

**personal email address required for electronic access.*

☐ Check enclosed ☐ Charge my (circle one): AMERICAN EXPRESS MASTER CARD VISA

CARD NUMBER _____ EXPIRATION DATE _____

NAME AND BILLING ADDRESS _____

SIGNATURE _____

Please fax this form to Natasha Felshman at 1 (212) 924-0240.

Payment must be included with this form. Your satisfaction is guaranteed or your money back.

If you wish to pay by check, please mail this form with payment to: EDventure Holdings, 104 Fifth Avenue, 20th Floor, New York, NY 10011, USA. If you have any questions, please call us at 1 (212) 924-8800; email us@edventure.com; www.edventure.com.

12-00